arXiv:1602.02452v1 [cs.CR] 8 Feb 2016

Privacy Preserving Architectures for Collaborative Intrusion Detection

Position Paper

Sashank Dara[1] and Dr.V.N Muralidhara[2]

[1] Cisco Systems Inc & IIIT-Bangalore
Bangalore, 560037, India
sadara@cisco.com
[2] IIIT-Bangalore
Bangalore, 560100, India
murali@iiitb.ac.in

Abstract. Collaboration among multiple organizations is imperative for contemporary intrusion detection. As modern threats become well sophisticated it is difficult for organizations to defend with threat context local to their networks alone. Availability of global *threat intelligence* is must for organizations to defend against modern advanced persistent threats (APTs). In order to benefit from such global context of attacks, privacy concerns continue to be of major hindrance. In this position paper we identify real world privacy problems as precise use cases, relevant cryptographic technologies and discuss privacy preserving architectures for collaborative intrusion detection.

1 Introduction

Contemporary intrusion detection is very daunting and challenging task due to two major factors. Firstly network infrastructures are becoming quite complex due to rapid technology advances[3]. Secondly modern cyber security attacks are very advanced, persistent and devastating. Threat actors are well funded, organized and such attacks happen at a global scale. Variety of information is needed to make modern intrusion detection possible and collaboration among different organizations is very imperative.

1.1 Telemetry

Identification of attacks would need *Telemetry*. This includes browsing history, network flows, file attributes, packet captures, firewall logs, application logs, system level behavior attributes etc. Verbose telemetry would facilitate event correlation, machine learning, security analytics and other techniques for identifying intrusions.

1.2 Threat Intelligence

Availability of telemetry will help in tracing specific attack artifacts like known bad domains, URLs serving malware, blacklisted IP addresses, malicious file signatures etc. Such precise attack artifacts are called *Indicators of Compromise(IOCs)* and are considered crucial for intrusion detection [1],[2].

[3] Due to areas like Cloud, BOYD, SDN/NFV, IOT etc.

*IOC*s together with other security data like IP reputations, user reputations, vulnerabilities, signatures etc are broadly termed as *Threat Intelligence(TI)*[3]. It is provided by various entities like commercial vendors, open source communities, government agencies, consortia etc.[4]

1.3 Privacy Challenges

Organizations would need to collaborate with their peers, threat intelligence providers and some times even other victims in order to defend against modern threats that happen at global scale[5],[6]. For such collaborative efforts to be successful it is essential for organizations to share their local telemetry[4]. Sharing telemetry would severely compromise the privacy of individuals and organizations. Often telemetry would contain

1. Sensitive *Personally Identifiable Information (PII)* like email-ids, user names, addresses, SSNs, credit card numbers, device identifiers like internal IP addresses etc.
2. *Protected Health Information (PHI)* like health related websites visited, health records, test reports etc.
3. Classified and confidential information like sensitive documents, trade secrets etc.

Many organizations are constrained to participate in such collaborative efforts due to privacy compliance, regional laws, risk of information leakage etc. and suffer from lacking global context of threats [7].

Contemporary privacy preserving methods are predominantly based on anonymization techniques, they have been well studied and determined to be prone to inferential attacks and other short comings [8],[9],[10],[11],[12].

1.4 Key Contributions

Although much of contemporary body of knowledge is about enhancing modern intrusion detection very little has been published on possible directions for privacy preservation. We believe leveraging the recent advances in cryptographic technologies would aide to better privacy preserving intrusion detection. In this position paper we identify real world problems as use cases, identify possible privacy technologies and architectures, we also discuss their threat models and challenges. We hope our work will motivate practical advances in this emerging field.

2 Prior Art

The immense benefit of collaborative distributed intrusion detection was first identified in [13]. Subsequent body of work witnessed major advancements in this area [14]. As the modern attacks became more sophisticated, advanced major research has been done in the areas of big data analytics, anomaly detection, machine learning etc applied for intrusion detection and emerged as very fast emerging field. It is very challenging to summarize these advances but good reference could be found here [15]. A comprehensive work on Collaborative Intrusion Detection is devoted to advantages of network of multiple intrusion detection systems in order to identify global threats[16]. Their work focuses very little on privacy issues. Also as we shall discuss in later sections there are many other forms of collaboration that are possible.

[4] More concrete use cases are provided in detail later

Application for *Secure Multiparty Computation* techniques for collaborative intrusion detection were first identified as open problem in [17]. Later much of the research progressed in privacy preserving data mining techniques[18]. But major break through work in this area of privacy preserving collaborative security monitoring appeared in [19], [20]. Subsequently work appeared in privacy preserving network outage monitoring [21], anomaly detection [22], traffic statistics [23], deep packet inspection[24], MiTB attacks [25], predictive blacklists[26][27], data leak detection[28], information aggregation [29], event correlation[30].

There is also exciting advances happening in cryptographic techniques collectively being called as *Computing on Encrypted Data(COED)*. It includes techniques like *Fully Homomorphic Encryption*, *Secure Multiparty Computation*, *Searchable Encryption* etc. A draft version describing these advances could be found here[31]. We identify relevant techniques needed for privacy preserving intrusion detection in next section.

3 Preliminaries

In this section we give informal introduction to different privacy preserving techniques in *COED* [5]. Some of these techniques are becoming mature while some are still in nascent stages.

3.1 Private Information Retrieval

Private Information Retrieval (PIR) techniques are an interesting cryptographic primitive. Informally in a typical *PIR* protocol, a *client* queries a *database server* privately such that the *server* returns the requested records without knowing what they are. As a trivial solution the server could return all the records and let the client chose desired ones locally but this has huge communication complexity. It is first shown in [32] that non-trivial *PIR* is impossible when only single database server exists in an information theoretic setting.

Information Theoretic (IT-PIR) protocols does not rely on any limitation of computing resources available with an adversary. These protocols are based on multiple database servers that not all of the servers are colluding with each other [33], [34], [35],[36]. *Computational PIR (CPIR)* are protocols rely on the computing resources available with an adversary for proving security guarantees [37], [38], [39], [40].

3.2 Secure Multiparty Computation

Secure Multiparty Computation(SMC) techniques enable multiple parties to compute some public function without revealing their individual inputs. Secure computations are carried out in a distributed manner on *secret shares* of the input data, the results are constructed back without ever revealing the inputs to participants. Recent advances in this technology are promising with open source libraries and commercial offerings emerging like SEPIA, SECREC, RMIND, FastGC, FairPlay, OblivC etc. A good overview of this could be found here [31].

[5] Formal definitions of the same are out of scope for this paper

4 Threat Intelligence Lookups

Many free and commercial services offer different types of *IOCs* as watch lists. Primarily there are two modes of operation, either the organizations query these online services on need basis or consume these watch lists as a feed for integrating into their internal analytics.

Subsequently they take appropriate action for identification of infections and remediation. Considering the sheer volume of such threat intelligence, many resource constrained organizations would prefer to do a lookup from their client machines to these online services than receiving them as feeds.

Looking up these Cloud based services needs sharing of *telemetry* which could contain private and sensitive information. Intrusion detection based on threat intelligence is predominantly about identification of *known-known*[6] cases in telemetry.

4.1 URL Reputation

Modern sophisticated attacks like *phishing, domain generated algorithms, command and control communication, data ex-filtration* are predominantly web based [41],[42]. In order to safe guard the users from malicious web based attacks, it is critical to know the reputation of the URL. Many prominent services exist that provide URL scanning service[43]. Given a particular URL, such services determine whether it is *known bad*. Sharing the web browsing data would severely breach the privacy of individual. In order to prevent such privacy breach, anonymization techniques proposed for URLs involve encrypting parts of the URL or truncating certain parameters etc. [44]. But such sanitization offers very poor privacy guarantees and are considered weak [8],[9],[10].

Other look up services are malicious public IP addresses[45], malware signatures [46] etc.

4.2 Architecture

Privacy preserving threat intelligence lookups is a classic case for *Private Information Retrieval(PIR)*. The users could keep their web browsing history (or network logs) private and lookup a public Cloud service in order to determine if they got infected. The proposed setup is shown in Figure.1

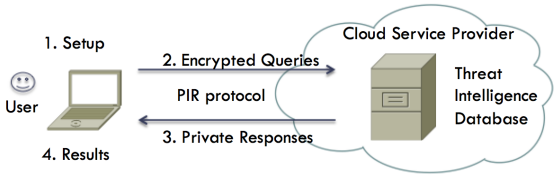

Fig. 1: PIR for Threat Intelligence Lookup

[6] Malware campaigns, Vulnerabilities that are publicly known

4.3 Threat Model

In this setup, privacy of *Users* is ensured from *Cloud Service Provider(CSP)*. The server is oblivious of the lookup strings and records that have been searched. Also *CPIR* protocols needs only one server in contrary to *IT-PIR* protocols that need many non colluding servers. In reality it is hard to prove if the servers of a single *CSP* are non-colluding so *CPIR* protocols are preferred.

4.4 Challenges

It has to be experimentally validated to determine how practical current schemes are. The schemes should be efficient many folds than trivial approach of downloading entire threat intelligence database to local network.

5 Collaborative Sharing and Analytics

5.1 Threat Intelligence Sharing

Many standards are emerging for sharing security related information [47]. Such standards enable organizations to share threat intelligence to prevent attacks in much earlier phases. Example recommendations that could be made with availability of collective attack observables could be like below

1. *40 out of 50 organizations similar to you are hit with this signature*
2. *28 out of 30 organizations in your geography has blocked his $port or $ip address or $domain*
3. *almost 98% of similar organizations have these top 20 rules as best practices ! "*

Currently projects like DSHIELD helps in identifying trends based on firewall logs. One of the major concerns still is privacy of information that could be shared in order for these efforts to be effective. Privacy preserving techniques that enable analysis on such attack observables are still in nascent stages.

5.2 Telemetry Sharing

Telemetry sharing and subsequent security analytics is an essential ingredient of modern intrusion detection for identifying *unknown* cases [48]. While analytics within the administrative domain only provide local context, global context like ongoing Distributed Denial of Service(DDoS) attacks, passive DNS analysis to identify botnets, Internet wide scans, unusual traffic patterns etc. would require analytics across multiple organizations.

5.3 Architecture

This architecture was first proposed for collaborative privacy preserving network security monitoring[19]. Network events were aggregated from different administrative domains and protocols were proposed to determine *top-k* nodes in a privacy preserving way. Sensitive network logs and events were *secret shared* and distributed among servers. Subsequently correlation were performed using *Secure Multi party Computation(SMC)* protocols such that no single server has access to entire inputs. Similar approach could be taken for collaborating on *Threat Intel Sharing* and *Telemetry Sharing* as well.

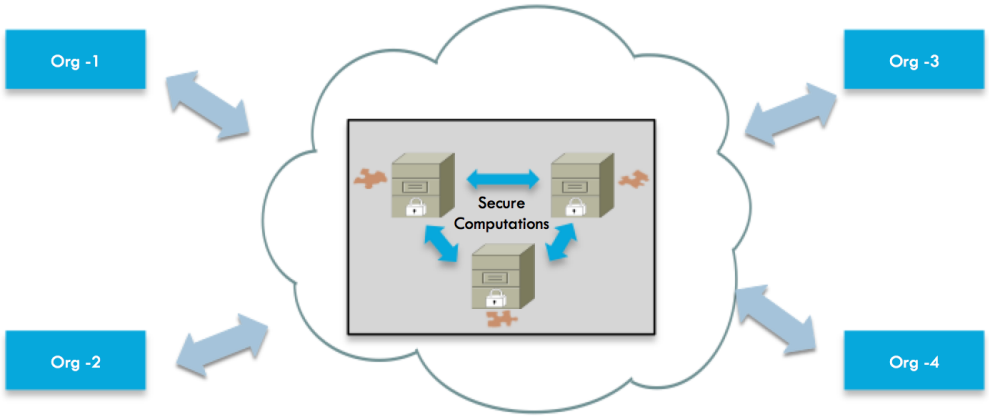

Fig. 2: Collaborative Security Analytics

5.4 Threat Model

In this setup none of the participants has access to private inputs of others. The results are made available to all participants as an incentive for participation. The distributed nodes performing *Secure Multiparty Computation (SMC)* can be run be consortium of organizations in a particular field or across fields. For example *Intelligence Sharing and Analysis Centers(ISAC)* of different fields like finance, banking, governments could run the infrastructure needed for *SMC* protocols. Subsequently the results to the participant organizations could be given as a value add in a privacy preserving way.

5.5 Challenges

Recent advances in tools like RMIND[49] are promising. They enable statistical analysis in a privacy preserving manner on *secret shares* of sensitive data to ensure their privacy. Rigorous experimentation of collaborative algorithms for traffic analysis, botnet detection etc. needs to be performed in order to evaluate their practical viability.

6 Conclusions

In this position paper we identified privacy preserving architectures for collaborative intrusion detection methods. Various challenges and advances needed in privacy preserving techniques are discussed throughout in realizing such architectures. As future work we intend to experimentally evaluate these proposed architectures in order to validate whether their feasibility in real world.

Acknowledgment

We would like to thank Cisco Systems for supporting this work.

References

1. http://www.21ct.com/default/assets/File/12-indicators-of-compromise.pdf.
2. https://www.sans.org/reading-room/whitepapers/detection/home-field-advantage-indicators-compromise-hunt-down-advanced-p
3. https://securosis.com/blog/applied-threat-intelligence-defining-ti.
4. https://www.alienvault.com/open-threat-exchange, https://www.facebook.com/threatexchange, http://www.threatconnect.com/platform/features/, http://www.talosintel.com/.
5. C. Gilbert, M. E. Hellman, and T. A. Berson, "Scalable security: Cyber threat information sharing," 2014.
6. K. Zetter, "Google hack attack was ultra sophisticated, new details show," *Wired Magazine*, vol. 14, 2010.
7. http://www.emc.com/collateral/industry-overview/11652-h9084-aptbdb-brf-0212-online.pdf.
8. R. Pang, M. Allman, V. Paxson, and J. Lee, "The devil and packet trace anonymization," *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 1, pp. 29–38, 2006.
9. S. E. Coull, C. V. Wright, A. D. Keromytis, F. Monrose, and M. K. Reiter, "Taming the devil: Techniques for evaluating anonymized network data," in *Network and Distributed System Security Symposium 2008: February 10-13, 2008, San Diego, California: Proceedings*. Internet Society, 2008, pp. 125–135.
10. J. King, K. Lakkaraju, and A. Slagell, "A taxonomy and adversarial model for attacks against network log anonymization," in *Proceedings of the 2009 ACM symposium on Applied Computing*. ACM, 2009, pp. 1286–1293.
11. T. Li and N. Li, "On the tradeoff between privacy and utility in data publishing," in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2009, pp. 517–526.
12. T. Brekne and A. Årnes, "Circumventing ip-address pseudonymization." in *Communications and Computer Networks*, 2005, pp. 43–48.
13. M. E. Locasto, J. J. Parekh, S. Stolfo, and V. Misra, "Collaborative distributed intrusion detection," 2004.
14. C. V. Zhou, C. Leckie, and S. Karunasekera, "A survey of coordinated attacks and collaborative intrusion detection," *Computers & Security*, vol. 29, no. 1, pp. 124–140, 2010.
15. S. Suthaharan, "Big data classification: Problems and challenges in network intrusion prediction with machine learning," *ACM SIGMETRICS Performance Evaluation Review*, vol. 41, no. 4, pp. 70–73, 2014.
16. C. Fung and R. Boutaba, *Intrusion Detection Networks: A Key to Collaborative Security*. CRC Press, 2013.
17. W. Du and M. J. Atallah, "Secure multi-party computation problems and their applications: a review and open problems," in *Proceedings of the 2001 workshop on New security paradigms*. ACM, 2001, pp. 13–22.
18. B. Fung, K. Wang, R. Chen, and P. S. Yu, "Privacy-preserving data publishing: A survey of recent developments," *ACM Computing Surveys (CSUR)*, vol. 42, no. 4, p. 14, 2010.
19. M. Burkhart, M. Strasser, D. Many, and X. Dimitropoulos, "Sepia: Privacy-preserving aggregation of multi-domain network events and statistics," *Network*, vol. 1, p. 101101, 2010.
20. M. Burkhart and X. Dimitropoulos, "Privacy-preserving distributed network troubleshootingbridging the gap between theory and practice," *ACM Transactions on Information and System Security (TISSEC)*, vol. 14, no. 4, p. 31, 2011.
21. M. Djatmiko, D. Schatzmann, X. Dimitropoulos, A. Friedman, and R. Boreli, "Collaborative network outage troubleshooting with secure multiparty computation," *Communications Magazine, IEEE*, vol. 51, no. 11, pp. 78–84, 2013.
22. P. Zhang, X. Huang, X. Sun, H. Wang, and Y. Ma, "Privacy-preserving anomaly detection across multi-domain networks," in *Fuzzy Systems and Knowledge Discovery (FSKD), 2012 9th International Conference on*. IEEE, 2012, pp. 1066–1070.

23. J. W. Brown, O. Ohrimenko, and R. Tamassia, "Haze: privacy-preserving real-time traffic statistics," in *Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. ACM, 2013, pp. 540–543.
24. J. Sherry, C. Lan, R. A. Popa, and S. Ratnasamy, "Blindbox: Deep packet inspection over encrypted traffic."
25. H. Qusa and S. Abudalfa, "Secure collaborative processing architecture for mitb attack detection," *International Journal of Network Security & Its Applications*, vol. 5, no. 5, p. 83, 2013.
26. L. Melis, A. Pyrgelis, and E. De Cristofaro, "Building and measuring privacy-preserving predictive blacklists," *arXiv preprint arXiv:1512.04114*, 2015.
27. J. Freudiger, E. De Cristofaro, and A. Brito, "Controlled data sharing for collaborative predictive blacklisting," in *Detection of Intrusions and Malware, and Vulnerability Assessment(DIMVA), 12th International Conference, DIMVA 2015, Milan, Italy,*. LNCS, 2015, pp. 327–349.
28. X. Shu and D. D. Yao, "Data leak detection as a service," in *Security and Privacy in Communication Networks*. Springer, 2013, pp. 222–240.
29. G. Kreitz, M. Dam, and D. Wikström, "Practical private information aggregation in large networks," in *Information Security Technology for Applications*. Springer, 2012, pp. 89–103.
30. H. Qusa, R. Baldoni, and R. Beraldi, "A privacy preserving scalable architecture for collaborative event correlation," in *Trust, Security and Privacy in Computing and Communications (TrustCom), 2012 IEEE 11th International Conference on*. IEEE, 2012, pp. 837–843.
31. https://www.cs.bris.ac.uk/~nigel/ECRYPT-MPC/Draft.pdf.
32. B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan, "Private information retrieval," *Journal of the ACM (JACM)*, vol. 45, no. 6, pp. 965–981, 1998.
33. A. Beimel, Y. Ishai, and T. Malkin, "Reducing the servers computation in private information retrieval: Pir with preprocessing," in *Advances in CryptologyCRYPTO 2000*. Springer, 2000, pp. 55–73.
34. I. Goldberg, "Improving the robustness of private information retrieval," in *Security and Privacy, 2007. SP'07. IEEE Symposium on*. IEEE, 2007, pp. 131–148.
35. R. Henry, F. Olumofin, and I. Goldberg, "Practical pir for electronic commerce," in *Proceedings of the 18th ACM conference on Computer and communications security*. ACM, 2011, pp. 677–690.
36. F. Olumofin and I. Goldberg, "Revisiting the computational practicality of private information retrieval," in *Financial Cryptography and Data Security*. Springer, 2012, pp. 158–172.
37. B. Chor and N. Gilboa, "Computationally private information retrieval," in *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*. ACM, 1997, pp. 304–313.
38. E. Kushilevitz and R. Ostrovsky, "Replication is not needed: Single database, computationally-private information retrieval," in *focs*. IEEE, 1997, p. 364.
39. Y.-C. Chang, "Single database private information retrieval with logarithmic communication," in *Information Security and Privacy*. Springer, 2004, pp. 50–61.
40. C. Aguilar-Melchor, J. Barrier, L. Fousse, and M.-O. Killijian, "Xpire: Private information retrieval for everyone," Cryptology ePrint Archive, Report 2014/1025, Tech. Rep., 2014.
41. E. M. Hutchins, M. J. Cloppert, and R. M. Amin, "Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains," *Leading Issues in Information Warfare & Security Research*, vol. 1, p. 80, 2011.
42. https://digital-forensics.sans.org/blog/2009/10/14/security-intelligence-attacking-the-kill-chain/.
43. http://www.urlblacklist.com/?sec=home, https://developers.google.com/safe-browsing/?hl=en, https://urlquery.net/, https://www.phishtank.com/.
44. G. Kuenning and E. L. Miller, "Anonymization techniques for urls and filenames," Technical report, University of California, Santa Cruz, Tech. Rep., 2003.
45. https://www.senderbase.org/, https://www.spamhaus.org/lookup/.
46. http://www.clamav.net/index.html, http://virusshare.com/, https://www.virustotal.com/.
47. http://www.mitre.org/capabilities/cybersecurity/overview/cybersecurity-resources/standards.
48. https://blogs.cisco.com/security/a-model-for-evaluating-breach-detection-readiness.

49. D. Bogdanov, L. Kamm, S. Laur, and V. Sokk, "Rmind: a tool for cryptographically secure statistical analysis." *IACR Cryptology ePrint Archive*, vol. 2014, p. 512, 2014.
50. https://www.snort.org/.
51. https://www.sandvine.com/downloads/general/global-internet-phenomena/2015/encrypted-internet-traffic.pdf.
52. D. Naylor, A. Finamore, I. Leontiadis, Y. Grunenberger, M. Mellia, M. Munafò, K. Papagiannaki, and P. Steenkiste, "The cost of the s in https," in *Proceedings of the 10th ACM International on Conference on emerging Networking Experiments and Technologies*. ACM, 2014, pp. 133–140.
53. https://tools.ietf.org/html/draft-mm-wg-effect-encrypt-01.
54. J. Jarmoc and D. S. C. T. Unit, "Ssl/tls interception proxies and transitive trust," *Black Hat Europe*, 2012.
55. L. S. Huang, A. Rice, E. Ellingsen, and C. Jackson, "Analyzing forged ssl certificates in the wild," in *Security and Privacy (SP), 2014 IEEE Symposium on*. IEEE, 2014, pp. 83–97.
56. https://tools.ietf.org/html/draft-hildebrand-middlebox-erosion-01.